

11-17-00

A

11/15/00
JC806 U.S. PTO

JCS71 U.S. PTO
09/714781
11/15/00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventorship..... Spelman et al.
Applicant..... Microsoft Corporation
Attorney's Docket No. MS1-649US
Title: Method and Apparatus for Generating Random Numbers

TRANSMITTAL LETTER AND CERTIFICATE OF MAILING

To: Commissioner of Patents and Trademarks,
Washington, D.C. 20231

From: Steven R. Sponseller (Tel. 509-324-9256; Fax 509-323-8979)
Lee & Hayes, PLLC
421 W. Riverside Avenue, Suite 500
Spokane, WA 99201

The following enumerated items accompany this transmittal letter and are being submitted for the matter identified in the above caption.

1. Specification—title page, plus 23 pages, including 37 claims and Abstract
2. Transmittal letter including Certificate of Express Mailing
3. 4 Sheets Formal Drawings (Figs. 1-4)
4. Return Post Card

Large Entity Status ☒ [x]

Small Entity Status ☐ []

Date: 11-15-2000

By: Steven R. Sponseller
Steven R. Sponseller
Reg. No. 39,384

CERTIFICATE OF MAILING

I hereby certify that the items listed above as enclosed are being deposited with the U.S. Postal Service as either first class mail, or Express Mail if the blank for Express Mail No. is completed below, in an envelope addressed to The Commissioner of Patents and Trademarks, Washington, D.C. 20231, on the below-indicated date. Any Express Mail No. has also been marked on the listed items.

Express Mail No. (if applicable) EL685271113

Date: 11-15-00

By: Lori A. Vierra
Lori A. Vierra

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Method and Apparatus for
Generating Random Numbers**

Inventors:

Scott A. Field

Jeffrey F. Spelman

ATTORNEY'S DOCKET NO. MS1-649US

1 **TECHNICAL FIELD**

2 The present invention relates to random number generators and, more
3 particularly, to computer-implemented random number generators that create
4 strings of random bits.

5
6 **BACKGROUND**

7 Random number generation is an important part of the security
8 infrastructure in many application programs and operating systems. For example,
9 random numbers are used to generate session keys and cryptographic keys for
10 encoding data that is transmitted between two locations (such as between a client
11 and a server). The use of such keys protects the integrity of the data and provides
12 for the authentication of the data and authentication of the user attempting to
13 access the data.

14 The quality of the random numbers generated is associated with the quality
15 of the security provided by the application program or operating system. A perfect
16 random number generator that produces a truly random sequence of bits is
17 considered by many to be impossible. Thus, designers attempt to create “pseudo”
18 random number generators that produce unpredictable sequences of bits in which
19 no particular bit is more likely to be generated at a given time or place in the
20 sequence than any other bit. This disclosure uses the terms “random number
21 generator” and “pseudo random number generator” interchangeably.

22 The quality of the random seed used by the random number generator
23 affects the quality of the random number created by the random number generator.
24 Common techniques for creating a random seed include using operating
25

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18
- 19
- 20
- 21
- 22
- 23
- 24
- 25

9
10
11
12
13
14
15
16
17
18

14
15
16
17
18

19

20
21
22
23
24
25

1 restarted. Further, the system and methods described herein collect entropy data
2 from multiple sources, thereby reducing the likelihood that two computer systems
3 will have the same entropy data. Thus, the systems and methods described herein
4 generate random numbers having an improved quality.

5 In one embodiment, entropy data is collected and stored in a nonvolatile
6 memory. The entropy data stored in the nonvolatile memory is updated with
7 newly collected entropy data. A string of random bits is generated from the
8 entropy data stored in the nonvolatile memory.

9 In a described embodiment, the entropy data is collected from multiple
10 sources within a computer system.

11 In a particular embodiment, the entropy data includes data related to a
12 processor in a computer system and data related to an operating system executing
13 on the computer system.

14 In a described implementation, the entropy data is maintained in a protected
15 portion of an operating system kernel such that the entropy data is inaccessible to
16 application programs executing on the system.

17 In one embodiment, generating a string of random bits includes hashing the
18 entropy data to generate random seed data.

19 A particular embodiment includes communicating the string of random bits
20 to an application program requesting a random number.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates a block diagram of a system that collects and stores entropy data and generates strings of random bits based on the entropy data.

Fig. 2 is a flow diagram illustrating a procedure for collecting, storing, and updating entropy data.

Fig. 3 is a flow diagram illustrating a procedure for generating a string of random bits at the request of an application program.

Fig. 4 illustrates an example of a suitable operating environment in which the random number generator may be implemented.

DETAILED DESCRIPTION

The system and methods described herein provide a random number generator that creates strings of random bits using entropy data that is collected over the lifetime of the computer system. Entropy data refers to the data (such as computer system state information) used as a random seed for the random number generator. Using data that is collected over the lifetime of the computer system improves the quality of the resulting random numbers because the data is not reset each time the computer system or application program is restarted. The collected data continues to change as additional operations, functions, and application programs are executed on the computer system. As time passes, the likelihood that two different computer systems would produce the same entropy data is reduced.

Fig. 1 illustrates a block diagram of a system 100 that collects and stores entropy data and generates strings of random bits based on the entropy data. In a

particular embodiment, the system 100 is contained within a computer system. The entropy data collected includes central processing unit (CPU) data 102 and operating system data 104. CPU data 102 includes various CPU state information and operating system data 104 includes various operating system state information. Table 1 below illustrates exemplary CPU data 102 and operating system data 104.

TABLE 1

CPU Data	Operating System Data
timestamp counter	boot time
cache misses per second	time of day
branch mispredictions per second	time zone bias
CPU-specific counters	page size
	number of processors
	current cache size
	peak cache size
	I/O read operation count
	I/O write operation count
	cache read count

The CPU data 102 may vary from one CPU to the next. For example, many of the internal counters in a CPU are affected by power fluctuations, the types of operations performed by the CPU, and the clock speed at which the CPU is operating.

In a particular embodiment, a system device driver resides in the operating system kernel and generates random numbers at the request of an application program or other function accessing the device driver. This device driver is responsible for collecting and maintaining entropy data as discussed herein. An application programming interface (API) is provided to allow application programs to request a random number. The API communicates random number requests to the device driver, which generates a random number (a string of random bits or bytes) based on the entropy data. In a particular implementation, application programs use the RtlGenRandom() API provided by the Windows[®]

operating system, developed by Microsoft Corporation of Redmond, Washington. The RtlGenRandom() API communicates with the device driver via the Win32[®] application programming interface call DeviceIoControl(), which is a commonly used API call for communicating with device drivers in the Windows[®] operating system.

As discussed below, the device driver applies a hash function to the various entropy data collected. The result of the hash function is used as the random seed for the random number generator. The entropy data and the random seed data are maintained in a protected portion of the operating system kernel (i.e., a portion of the operating system kernel that is not accessible by an application program). Maintaining the entropy data and the random seed data in a protected portion of the operating system kernel prevents an application program from predicting or deriving random numbers issued to another application program on the same computer system. In a particular embodiment, the device driver manages the memory used to store the entropy data and the random seed data. In this embodiment, entropy data is maintained in the operating system kernel as well as the non-volatile Windows[®] registry.

A typical computer system has multiple processes executing simultaneously, one or more of which may require random numbers. The system and methods described herein allow the generation of multiple random numbers for use as session keys, cryptographic keys, and the like. Although particular embodiments are discussed with reference to a device driver residing in the operating system kernel that generates random numbers, it will be appreciated that any type of software component and/or firmware component can be used to implement the random number generator.

Fig. 2 is a flow diagram illustrating a procedure 200 for collecting, storing, and updating entropy data. When a system is first initialized, a device driver collects initial entropy data from the CPU and the operating system (block 202). The initial entropy data is then stored in a nonvolatile memory (block 204), such as nonvolatile memory 108 in Fig. 1. The initial entropy data is hashed to generate random seed data for the random number generator (block 206). The hashed data may be stored in a register or other storage location that is accessible to the random number generator, but inaccessible to application programs executing on the system. Any hashing algorithm that produces a long string of bits can be used to hash the entropy data. An example hashing algorithm is described in U.S. Patent 5,778,069, the disclosure of which is incorporated by reference herein. In an alternate embodiment, two or more different hashing algorithms are applied to the same set of entropy data and the results are concatenated together into a single string of bits representing the random seed data.

At block 208 in Fig. 2, the procedure 200 determines whether to update the entropy data. The entropy data may be updated at periodic intervals (e.g., every fifteen minutes) or after generating a particular number of random numbers (e.g., after every tenth random number is generated). Alternatively, an application program may specifically request an update of the entropy data. A particular implementation updates the data the first time that an application program makes a request for a random number. If the entropy data needs to be updated, then the procedure continues to block 210, where the device driver collects the current entropy data (i.e., the CPU data and the operating system data). After collecting the current entropy data, the device driver updates the data in the nonvolatile

1	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	10
11	11
12	12
13	13
14	14
15	15
16	16
17	17
18	18
19	19
20	20
21	21
22	22
23	23
24	24
25	25
26	26
27	27
28	28
29	29
30	30
31	31
32	32
33	33
34	34
35	35
36	36
37	37
38	38
39	39
40	40
41	41
42	42
43	43
44	44
45	45
46	46
47	47
48	48
49	49
50	50
51	51
52	52
53	53
54	54
55	55
56	56
57	57
58	58
59	59
60	60
61	61
62	62
63	63
64	64
65	65
66	66
67	67
68	68
69	69
70	70
71	71
72	72
73	73
74	74
75	75
76	76
77	77
78	78
79	79
80	80
81	81
82	82
83	83
84	84
85	85
86	86
87	87
88	88
89	89
90	90
91	91
92	92
93	93
94	94
95	95
96	96
97	97
98	98
99	99
100	100

9
10
11
12
13
14
15
16
17

18
19
20
21
22
23
24
25

In a particular embodiment, the entropy data is hashed to produce a 640 bit hash, which is the seed data for the random number generator. The random number generator uses the 640 bit hash to generate a 256 byte random number, which is also referred to as a “key.” The 256 byte random number can be used as a session key, a cryptographic key, or in any other situation requiring a random number. In one implementation, the RSA RC4 stream cipher (available from RSA Security of Bedford, Massachusetts) is used to generate a 256 byte random number from the 640 bit hash.

Although particular implementations have been described above with reference to specific stream ciphers, other types of ciphers can be used to generate a random number from the 640 bit hash. Further, the 640 bit hash and the 256 byte random numbers represent an exemplary embodiment. The system and methods described herein can be used with a hash (i.e., seed data) of any size to generate a random number having any number of bits (or bytes).

Fig. 4 illustrates an example of a suitable operating environment in which the random number generator may be implemented. The illustrated operating environment is only one example of a suitable operating environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Other well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, gaming consoles, cellular telephones, network PCs, minicomputers,

1 mainframe computers, distributed computing environments that include any of the
2 above systems or devices, and the like.

3 Fig. 4 shows a general example of a computer 342 that can be used in
4 accordance with the invention. Computer 342 is shown as an example of a
5 computer that can perform the hashing and random number generation functions
6 described herein. Computer 342 includes one or more processors or processing
7 units 344, a system memory 346, and a bus 348 that couples various system
8 components including the system memory 346 to processors 344.

9 The bus 348 represents one or more of any of several types of bus
10 structures, including a memory bus or memory controller, a peripheral bus, an
11 accelerated graphics port, and a processor or local bus using any of a variety of
12 bus architectures. The system memory 346 includes read only memory (ROM)
13 350 and random access memory (RAM) 352. A basic input/output system (BIOS)
14 354, containing the basic routines that help to transfer information between
15 elements within computer 342, such as during start-up, is stored in ROM 350.
16 Computer 342 further includes a hard disk drive 356 for reading from and writing
17 to a hard disk, not shown, connected to bus 348 via a hard disk drive interface 357
18 (e.g., a SCSI, ATA, or other type of interface); a magnetic disk drive 358 for
19 reading from and writing to a removable magnetic disk 360, connected to bus 348
20 via a magnetic disk drive interface 361; and an optical disk drive 362 for reading
21 from and/or writing to a removable optical disk 364 such as a CD ROM, DVD, or
22 other optical media, connected to bus 348 via an optical drive interface 365. The
23 drives and their associated computer-readable media provide nonvolatile storage
24 of computer readable instructions, data structures, program modules and other data
25 for computer 342. Although the exemplary environment described herein employs

1 a hard disk, a removable magnetic disk 360 and a removable optical disk 364, it
2 will be appreciated by those skilled in the art that other types of computer readable
3 media which can store data that is accessible by a computer, such as magnetic
4 cassettes, flash memory cards, random access memories (RAMs), read only
5 memories (ROM), and the like, may also be used in the exemplary operating
6 environment.

7 A number of program modules may be stored on the hard disk, magnetic
8 disk 360, optical disk 364, ROM 350, or RAM 352, including an operating system
9 370, one or more application programs 372, other program modules 374, and
10 program data 376. A user may enter commands and information into computer
11 342 through input devices such as keyboard 378 and pointing device 380. Other
12 input devices (not shown) may include a microphone, joystick, game pad, satellite
13 dish, scanner, or the like. These and other input devices are connected to the
14 processing unit 344 through an interface 368 that is coupled to the system bus
15 (e.g., a serial port interface, a parallel port interface, a universal serial bus (USB)
16 interface, etc.). A monitor 384 or other type of display device is also connected to
17 the system bus 348 via an interface, such as a video adapter 386. In addition to the
18 monitor, personal computers typically include other peripheral output devices (not
19 shown) such as speakers and printers.

20 Computer 342 operates in a networked environment using logical
21 connections to one or more remote computers, such as a remote computer 388.
22 The remote computer 388 may be another personal computer, a server, a router, a
23 network PC, a peer device or other common network node, and typically includes
24 many or all of the elements described above relative to computer 342, although
25 only a memory storage device 390 has been illustrated in Fig. 4. The logical

connections depicted in Fig. 4 include a local area network (LAN) 392 and a wide area network (WAN) 394. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet. In certain embodiments, computer 342 executes an Internet Web browser program (which may optionally be integrated into the operating system 370) such as the “Internet Explorer” Web browser manufactured and distributed by Microsoft Corporation of Redmond, Washington.

When used in a LAN networking environment, computer 342 is connected to the local network 392 through a network interface or adapter 396. When used in a WAN networking environment, computer 342 typically includes a modem 398 or other means for establishing communications over the wide area network 394, such as the Internet. The modem 398, which may be internal or external, is connected to the system bus 348 via a serial port interface 368. In a networked environment, program modules depicted relative to the personal computer 342, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

Computer 342 typically includes at least some form of computer readable media. Computer readable media can be any available media that can be accessed by computer 342. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to,

1 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
2 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
3 tape, magnetic disk storage or other magnetic storage devices, or any other media
4 which can be used to store the desired information and which can be accessed by
5 computer 342. Communication media typically embodies computer readable
6 instructions, data structures, program modules or other data in a modulated data
7 signal such as a carrier wave or other transport mechanism and includes any
8 information delivery media. The term "modulated data signal" means a signal that
9 has one or more of its characteristics set or changed in such a manner as to encode
10 information in the signal. By way of example, and not limitation, communication
11 media includes wired media such as wired network or direct-wired connection,
12 and wireless media such as acoustic, RF, infrared and other wireless media.
13 Combinations of any of the above should also be included within the scope of
14 computer readable media.

15 The invention has been described in part in the general context of
16 computer-executable instructions, such as program modules, executed by one or
17 more computers or other devices. Generally, program modules include routines,
18 programs, objects, components, data structures, etc. that perform particular tasks
19 or implement particular abstract data types. Typically the functionality of the
20 program modules may be combined or distributed as desired in various
21 embodiments.

22 For purposes of illustration, programs and other executable program
23 components such as the operating system are illustrated herein as discrete blocks,
24 although it is recognized that such programs and components reside at various
25

1 times in different storage components of the computer, and are executed by the
2 data processor(s) of the computer.

3 Thus, a system and method has been described that generate random
4 numbers based on entropy data collected over the lifetime of the computer system.
5 The entropy data is maintained in a persistent storage device and can be updated at
6 regular intervals.

7 Although the description above uses language that is specific to structural
8 features and/or methodological acts, it is to be understood that the invention
9 defined in the appended claims is not limited to the specific features or acts
10 described. Rather, the specific features and acts are disclosed as exemplary forms
11 of implementing the invention.
12
13
14
15
16
17
18
19
20
21
22
23
24
25

1 **CLAIMS**

2 1. A method comprising:
3 collecting entropy data;
4 storing the entropy data in a nonvolatile memory;
5 updating the entropy data stored in the nonvolatile memory with newly
6 collected entropy data; and
7 generating a string of random bits from the entropy data stored in the
8 nonvolatile memory.

9
10 2. A method as recited in claim 1 wherein the entropy data is collected
11 from multiple sources.

12
13 3. A method as recited in claim 1 wherein the entropy data is collected
14 from multiple sources within a computer system.

15
16 4. A method as recited in claim 1 wherein the entropy data includes data
17 related to a processor in a computer system.

18
19 5. A method as recited in claim 1 wherein the entropy data includes data
20 related to an operating system executing on a computer system.

21
22 6. A method as recited in claim 1 wherein the entropy data is maintained
23 in a protected portion of an operating system kernel.
24
25

1 7. A method as recited in claim 1 wherein the method is executing on a
2 system and the entropy data is inaccessible by an application program executing
3 on the system.

4
5 8. A method as recited in claim 1 wherein generating a string of random
6 bits includes hashing the entropy data to generate random seed data.

7
8 9. A method as recited in claim 1 wherein updating the entropy data
9 stored in the nonvolatile memory includes collecting new entropy data at periodic
10 intervals.

11
12 10. A method as recited in claim 1 further including communicating the
13 string of random bits to an application program requesting a random number.

14
15 11. One or more computer-readable memories containing a computer
16 program that is executable by a processor to perform the method recited in claim
17 1.

18
19 12. A method comprising:
20 receiving a request for a random number;
21 retrieving entropy data from a nonvolatile memory device, wherein the
22 entropy data is regularly updated with newly collected entropy data;
23 hashing the entropy data to create random seed data;
24 generating a string of random bits from the random seed data; and
25

communicating the string of random bits to the requester of the random number.

13. A method as recited in claim 12 wherein the entropy data is collected from multiple sources within a computer system.

14. A method as recited in claim 12 wherein the entropy data includes data related to a state of a processor in a computer system and data related to a state of an operating system executing on the computer system.

15. A method as recited in claim 12 wherein the entropy data is maintained in a protected portion of an operating system kernel.

16. A method as recited in claim 12 wherein the random seed data is maintained in a protected portion of an operating system kernel.

17. A method as recited in claim 12 wherein the entropy data is inaccessible by the requester of the random number.

18. One or more computer-readable memories containing a computer program that is executable by a processor to perform the method recited in claim 12.

1 **19.** A method comprising:
2 collecting entropy data;
3 storing the entropy data in a protected portion of an operating system
4 kernel; and
5 generating a string of random bits based on the entropy data.

6
7 **20.** A method as recited in claim 19 wherein the entropy data is
8 collected from multiple sources.

9
10 **21.** A method as recited in claim 19 wherein the entropy data is
11 inaccessible by an application program.

12
13 **22.** A method as recited in claim 19 further comprising updating the
14 entropy data with newly collected entropy data.

15
16 **23.** A method as recited in claim 19 further comprising communicating
17 the string of random bits to an application program requesting a random number.

18
19 **24.** One or more computer-readable memories containing a computer
20 program that is executable by a processor to perform the method recited in claim
21 19.

1 **25.** An apparatus comprising:

2 a nonvolatile memory configured to store entropy data, wherein the entropy
3 data stored in the nonvolatile memory is updated regularly; and

4 a random number generator coupled to the nonvolatile memory, wherein
5 the random number generator utilizes the entropy data stored in the nonvolatile
6 memory to generate strings of random bits.

7
8 **26.** An apparatus as recited in claim 25 wherein the entropy data is
9 collected from multiple sources.

10
11 **27.** An apparatus as recited in claim 25 wherein the entropy data is
12 updated at periodic intervals.

13
14 **28.** An apparatus as recited in claim 25 wherein the entropy data is
15 maintained in a protected portion of an operating system kernel such that the
16 entropy data is inaccessible by an application program.

17
18 **29.** An apparatus as recited in claim 25 wherein the entropy data
19 includes data related to a processor in a computer system and an operating system
20 executing on the computer system.

21
22 **30.** An apparatus as recited in claim 25 wherein the random number
23 generator hashes the entropy data to generate random seed data.

1 **31.** An apparatus as recited in claim 25 further including a timer
2 coupled to the random number generator, the timer indicating when to update the
3 entropy data stored in the nonvolatile memory device.

4
5 **32.** One or more computer-readable media having stored thereon a
6 computer program that, when executed by one or more processors, causes the one
7 or more processors to:

8 collect entropy data from multiple sources;
9 store the collected entropy data in a nonvolatile memory;
10 update the entropy data stored in the nonvolatile memory with newly
11 collected entropy data; and
12 produce a string of random bits from the entropy data stored in the
13 nonvolatile memory.

14
15 **33.** One or more computer-readable media as recited in claim 32
16 wherein the entropy data includes data related to a state of one or more processors.

17
18 **34.** One or more computer-readable media as recited in claim 32
19 wherein the entropy data is maintained in a protected portion of an operating
20 system kernel.

21
22 **35.** One or more computer-readable media as recited in claim 32
23 wherein the entropy data includes data related to a state of an operating system
24 executing on a computer system.

25

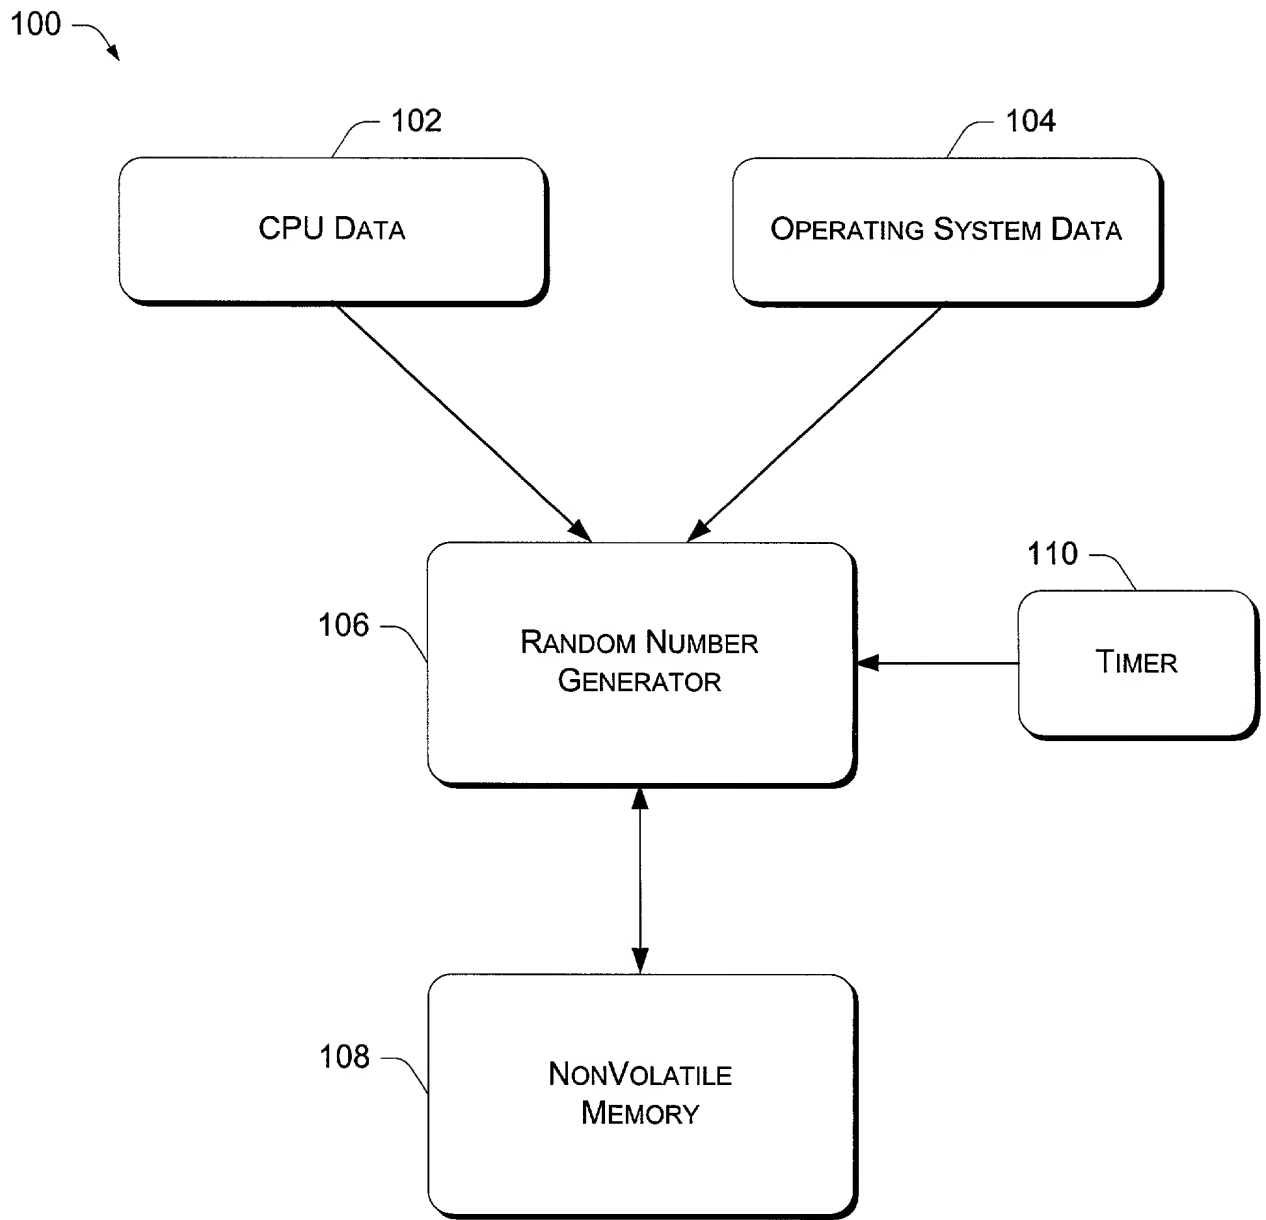
1 **36.** One or more computer-readable media as recited in claim 32
2 wherein to produce a string of random bits from the entropy data, the one or more
3 processors hash the entropy data to generate random seed data.

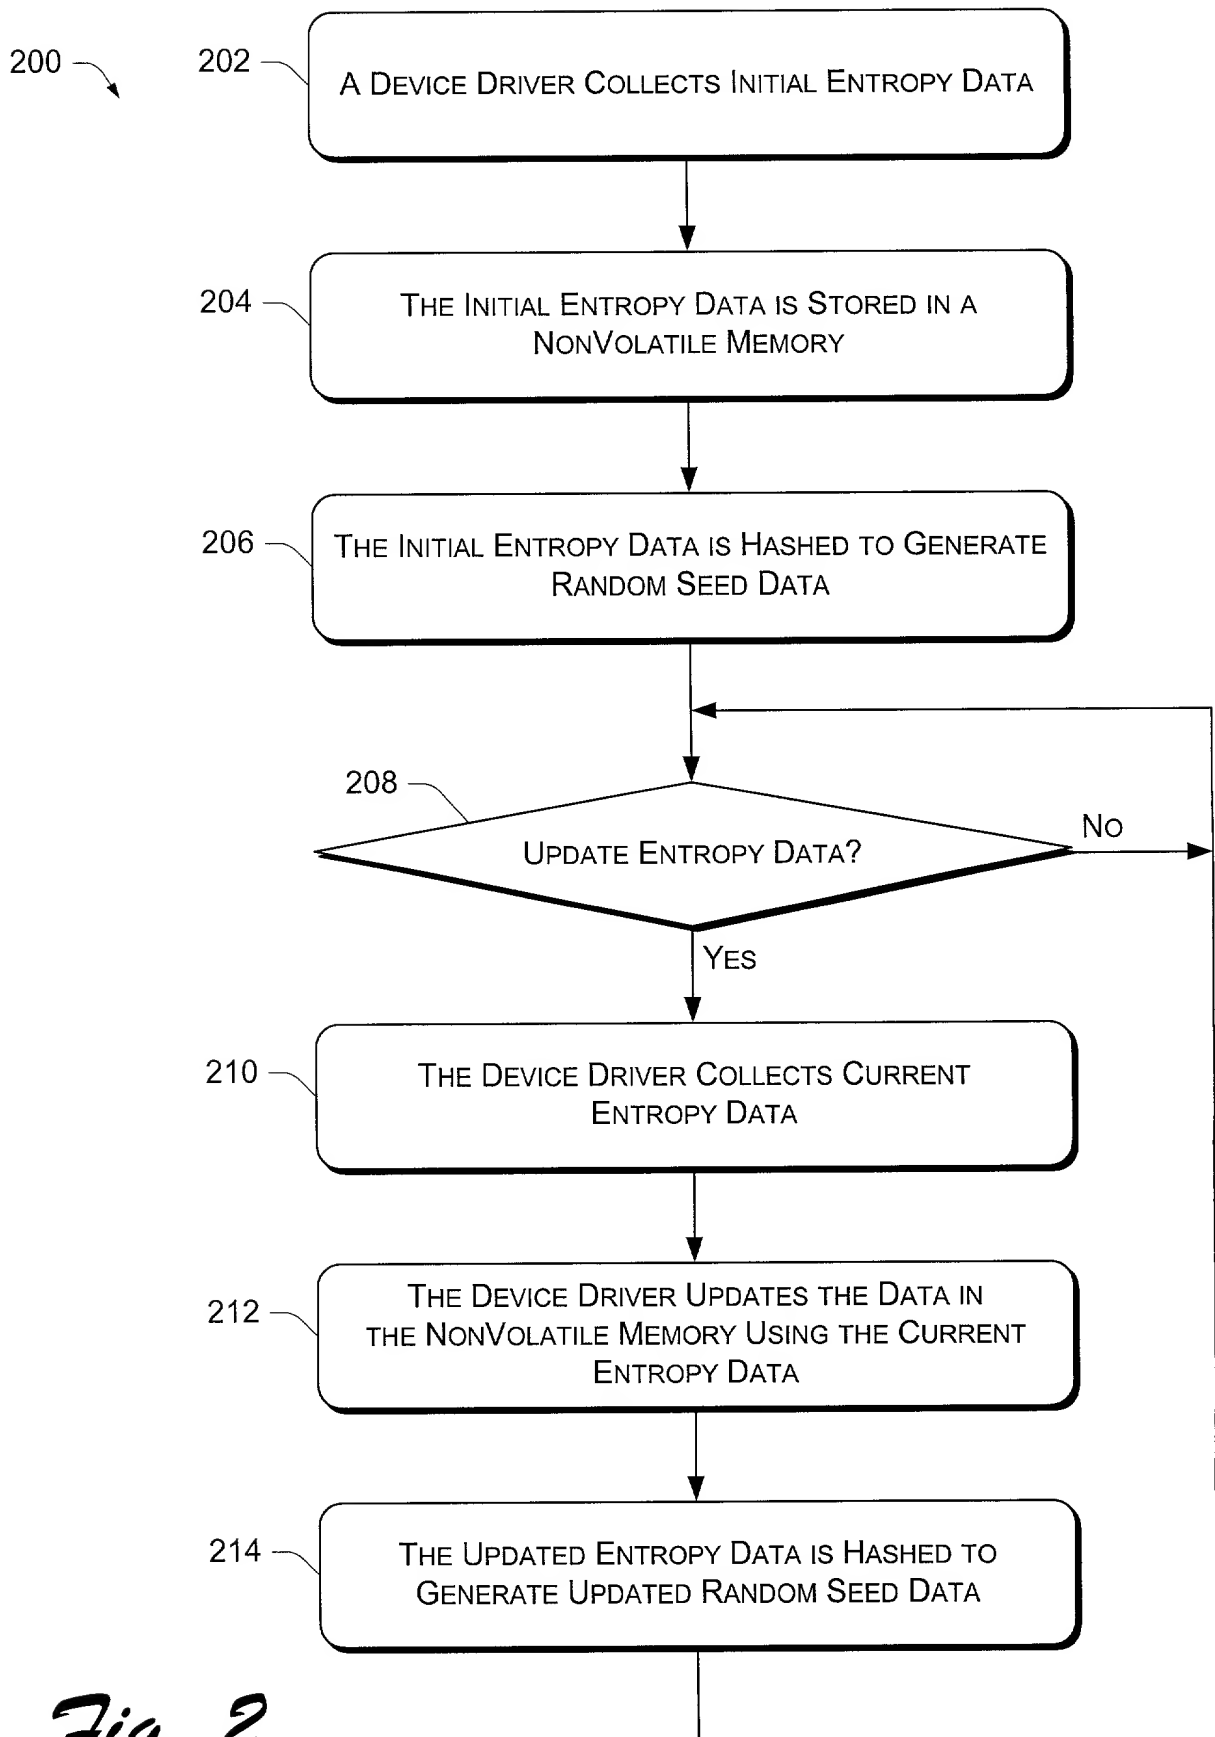
4
5 **37.** One or more computer-readable media as recited in claim 32
6 wherein the entropy data stored in the nonvolatile memory is updated with newly
7 collected entropy data at periodic intervals.

8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

1 **ABSTRACT**

2 A system collects entropy data and stores the entropy data in a nonvolatile
3 memory. The entropy data stored in the nonvolatile memory is updated with
4 newly collected entropy data. The entropy data stored in the nonvolatile memory
5 is used to generate a string of random bits. The entropy data is collected from
6 multiple sources within a computer system and may include data related to a
7 processor in the computer system and an operating system executing on the
8 computer system. The entropy data is maintained in a protected portion of an
9 operating system kernel. A hashing algorithm is applied to the entropy data to
10 generate random seed data.
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

*Fig. 1*

*Fig. 2*

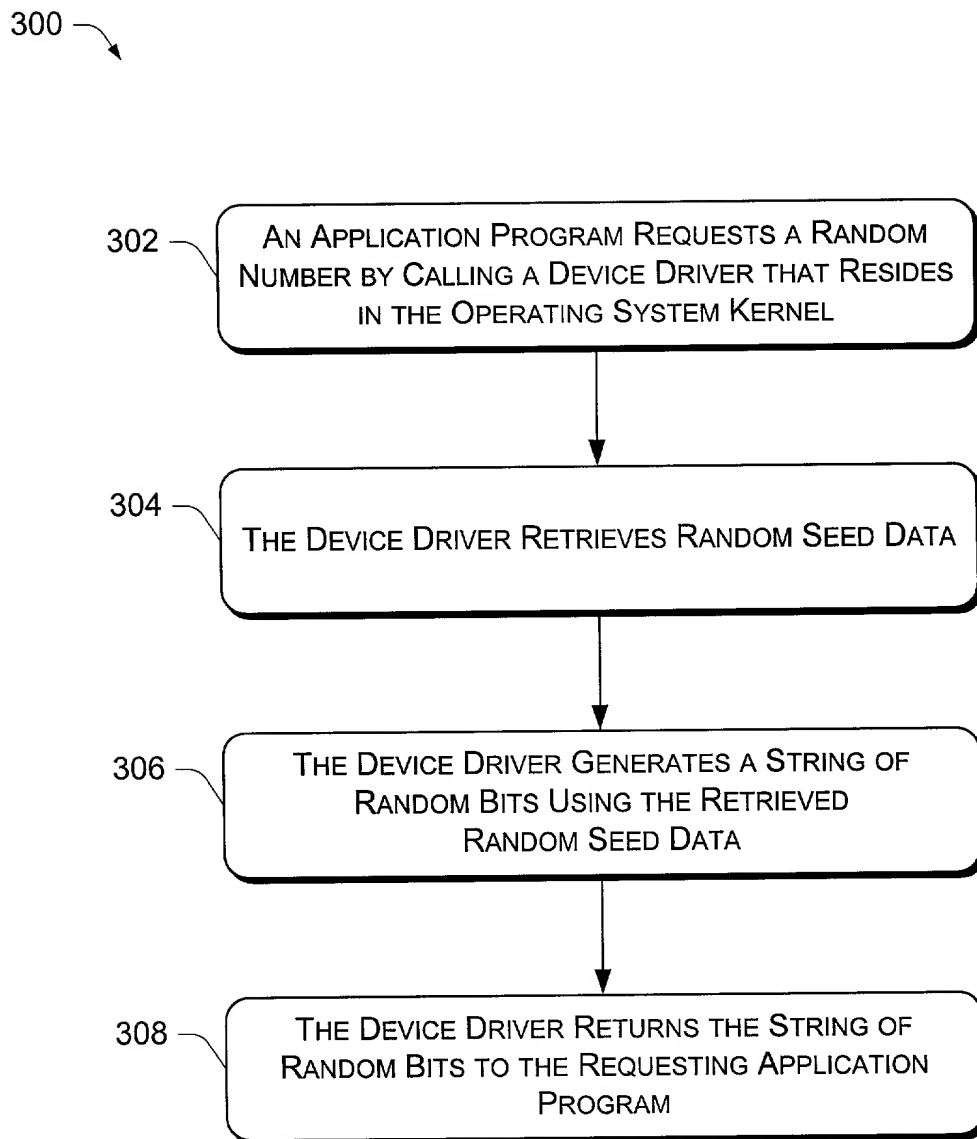
*Fig. 3*

Fig. 4

